arXiv:1512.04086v3 [cs.CV] 16 Feb 2016

Deep Learning-Based Image Kernel for Inductive Transfer

Neeraj Kumar, Animesh Karmakar, Ranti Dev Sharma, Abhinav Mittal, Amit Sethi
(NEERAJ.KUMAR,A.KARMAKAR,ABHINAV.MITTAL,RANTI,AMITSETHI)@IITG.ERNET.IN
Indian Institute of Technology Guwahati, Assam, India, 780139

Abstract

We propose a method to classify images from target classes with a small number of training examples based on transfer learning from non-target classes. Without using any more information than class labels for samples from non-target classes, we train a Siamese net to estimate the probability of two images to belong to the same class. With some post-processing, output of the Siamese net can be used to form a gram matrix of a Mercer kernel. Coupled with a support vector machine (SVM), such a kernel gave reasonable classification accuracy on target classes without any fine-tuning. When the Siamese net was only partially fine-tuned using a small number of samples from the target classes, the resulting classifier outperformed the state-of-the-art and other alternatives. We share class separation capabilities and insights into the learning process of such a kernel on MNIST, Dogs vs. Cats, and CIFAR-10 datasets.

1. Introduction

Deep learning architectures, notably the variants of convolutional neural networks (CNNs), have produced state-of-the-art results in large as well as very large-scale image classification and recognition problems (Hinton et al., 2012; Simonyan & Zisserman, 2014). The success of CNNs can be attributed to their capabilities of learning a hierarchy of increasingly complex and class-specific features using efficient training algorithms with appropriate connectivity and weight-sharing constraints in the convolutional layers. However, the practical utility of CNNs is limited in situations where training data is limited due to their large training sample requirement to achieve acceptable recognition rates. For instance, some of the top performing variant of CNN on CIFAR-10 dataset use 50,000 training samples and take a few days to train on a single workstation (Hinton et al., 2012). In this paper, we focus on reducing the training sample and time requirement of CNN-based techniques while still hoping to utilize their hierarchical feature learning capabilities for high classification accuracy.

Before the explosion of deep learning, use of support vector machines (SVMs) with various hand-crafted features represented the state-of-the-art for image recognition (Chapelle et al., 1999). Although, SVM-based systems produced lower peak recognition rates, their main advantage was lower number of training sample and time requirement than their deep learning counterparts (Hinton et al., 2012). SVMs have been coupled with variants of CNN by replacing the latter's classification layer (usually a softmax layer). In some cases, SVM not only functions as a wide margin classifier, it also provides cost and gradient values for CNN training. But due to gradient dilution at the lower layers of CNN, a large number of training samples and training time are still required for good results (Hinton et al., 2012). It has also been shown that CNN features[1] that was pre-trained in a supervised manner on non-target classes[2] can be used as an image representation in an SVM for target classes to give surprisingly good recognition performance (Huang & LeCun, 2006b). Such transfer learning from non-target classes is our main inspiration for experimenting with a different architecture that improves upon their recognition rates for a small number of training samples from target classes.

We propose an SVM-based classifier that operates on a trainable kernel based on a Siamese deep neural network (henceforth, Siamese net)[3] (Chopra et al., 2005). Assuming lack of a notion of semantic similarity, we propose that

Proceedings of the 33^{rd} International Conference on Machine Learning, New York, NY, USA, 2016. JMLR: W&CP volume 48. Copyright 2016 by the author(s).

[1] CNN features is a term that usually means a flattened vector of the output of the final convolutional layer of a CNN.

[2] We assume that we are interested in classifying samples from a *target* set of classes with a small number of labeled samples, while we have access to a large number of labeled examples from separate set of *non-target* classes for transfer learning.

[3] A Siamese network usually has an identical pair of convolutional and pooling layer stacks that share a single stack of fully connected layers on top. It is mainly trained and used to compute similarity between the pair of input images.

the Siamese net can be trained to estimate the probability that its input image pair belongs to the same class, *no matter what that class is*. Our hope is to be able to apply such Siamese net to target classes with no to little fine-tuning, which requires that its learning of similarity using non-target classes generalize to target classes.

Although many transfer learning methods for CNNs also use only the class labels for pre-training on non-target classes, the hope in CNNs is to learn discriminative features that generalize from non-target classes to target classes. We compare these two learning approaches in terms of their image recognition performance on standard datasets – MNIST digits (MNIST), Dogs vs. Cats (Dogs vs. Cats), and CIFAR-10 objects (Alex Krizhevsky & Hinton). We also experimented with two different levels of transfer learning – one in which the fully connected layers on top of the Siamese convolutional stacks are fine-tuned on target classes, and another in which even the fully connected layers are trained on non-target classes. In both cases, we reduced the practical training time by not training the convolutional layer at all, and copying weights from CNNs trained by others. Yet, in most cases, our kernels outperformed static kernels (RBF and linear) operating on CNN features.

For use in an SVM, the output of the Siamese net is not guaranteed to be positive semi-definite (PSD), which is one of the Mercers conditions (Schölkopf & Smola, 2002). To remedy this, we propose to use a transformation of the gram matrix produced by the Siamese net to a PSD matrix. We refer to this kernel as DEep Semantic Kernel (DESK).

2. Related Work

CNNs, starting from LeNet (LeNet), ushered the popularity of deep learning by improving object recognition rates significantly over previous approaches. Their main benefit is that they do not rely on hand-crafted features and learn an increasingly complex feature hierarchy from the data itself. To improve the image recognition performance of CNNs, various changes in hyperparameters and the mathematical functions in each layer and their training algorithms have been proposed (Boureau et al., 2010; Dahl et al., 2013; Srivastava et al., 2014). It has also been demonstrated that a linear SVM used in the final layer instead of logistic sigmoid for binary classification may improve the recognition performance (Huang & LeCun, 2006a). However, the number of training samples required for even these variants of CNN remains large because they retain the need to train a deep stack of layers using gradient descent, where the gradient dilutes away from the output layer (Hinton et al., 2012).

Use of unlabeled or labeled data from non-target classes in a transfer learning setting has become a common practice in CNNs (Huang & LeCun, 2006b; Philip Bachman, 2014; Alexey Dosovitskiy, 2015). Practical ability to utilize transfer learning has increased with the advent of deep learning libraries such as Theano (Theano), Caffe (Caffe), and researchers' willingness to make their models trained using these libraries accessible online. Transfer learning is done mainly by simply copying or initializing the weights of initial layers that are most affected by gradient dilution using models trained on other classes. It has been shown that the first few layers learn class-agnostic features that are most amenable to transfer learning (Yosinski et al., 2014). An SVM trained on CNN features obtained using non-target classes gives decent classification performance out of the box (Yosinski et al., 2014). Other approaches to transfer learning in CNNs besides using pre-trained convolutional layers include the use of successive frames of a video for unsupervised pre-training (Zou et al., 2011). However, for state-of-the-art performance layers are still fine-tuned based on tens, if not hundreds, of thousands of training examples from target classes, which also reflects in their training times.

With DESK, we take the approach of learning a kernel instead of features for use in SVM using Siamese nets – a variant of CNN. A Siamese net, such as the one used in DESK, computes similarity s_{ij} between images i and j. This approach has been used in recent works to learn distance between faces for verification (Junlin et al., 2014) and similarity learning between image patches for wide baseline stereo matching (Zagoruyko & Komodakis, 2015). With DESK, we went beyond similarity learning to propose a complete kernel-based classification framework that also utilizes transfer learning to yield high classification accuracy with a small training sample size.

Kernel learning outside of deep learning has been an active area of research. Most notable successes have been multiple kernel learning (Sonnenburg et al., 2006), (Gönen & Alpaydin, 2011), hierarchical arc-cosine kernels (Cho & Saul, 2009), and use of semi-definite programming to learn a kernel matrix (Lanckriet et al., 2004). While these works have shown improvement over the use of static kernels for image recognition, due to their reliance on hand-crafted and shallow features, these have been outperformed by their deep learning counterparts with the exception of (Cho & Saul, 2009).

Attempt to show equivalence of hierarchical kernels and deep learning include (Montavon et al., 2011) and (Cho & Saul, 2009), while use of multiple kernels in a convolutional architecture was proposed by (Mairal et al., 2014). These works further the theoretical understanding of deep learning. Our attempt is in line with these efforts but is aimed at reducing the training time and sample requirement

of CNNs without compromising on object recognition accuracy.

3. DESK Architecture and Training

With DESK, our goals in addition to learning similarity between images from target classes were the following:

1. *Propose an efficient method* to learn to compute similarity between pairs of images from target classes in terms of sample and time requirement.
2. *Study trade-off* between the extent of transfer learning from a kernel trained on non-target classes and target classes and classification accuracy.
3. *Probe for differences* between learning a deep kernel and deep features for classification.
4. *Convert the similarity score into a Mercer kernel* to train an SVM for classification of target classes.

DESK's Siamese net architecture takes two images to be compared as inputs and computes their similarity. This similarity score is converted into a Mercer kernel using a post-processing step. The kernel is then used to train an SVM using input images from target classes. This is shown in Figure 1. We start with presenting different options for DESK's neural network (NN).

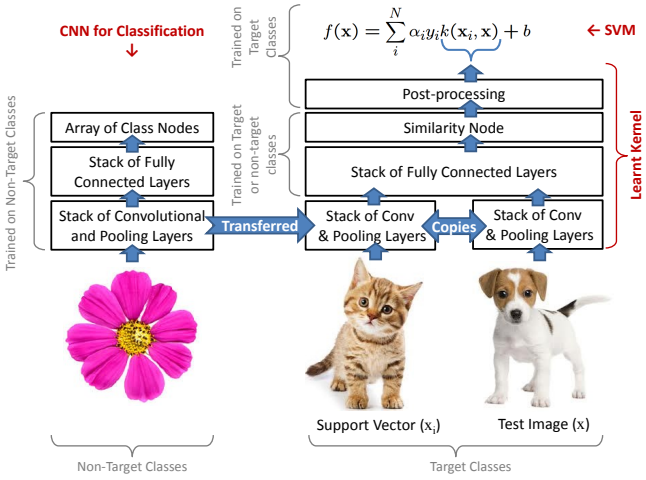

Figure 1. Proposed classification scheme based on an SVM using a trainable kernel that lends itself to transfer learning.

3.1. NN architecture

The defining features of the architectures that we experimented with were the following:

1. *Input*: Two images, i and j, to be compared for similarity were taken as input.

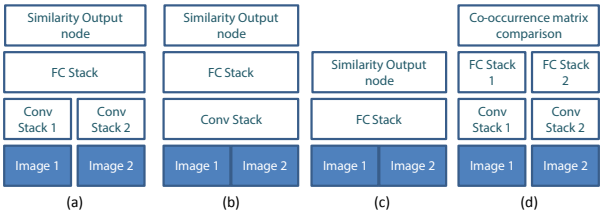

Figure 2. Various possible architectures for the neural network of the trainable kernel: (a) Siamese, (b) Super-Image, (c) Only fully-connected layer, (d) Deep-feature co-occurrence matrix.

2. *Output*: A target output for supervised training that was a scalar (similarity score), unlike a vector in case of multi-class CNNs.
3. *Convolutional layers*: A stack of convolutional layers (includes convolution, nonlinear squashing, and pooling) extracted increasingly complex features from the input images.
4. *Fully connected layers*: A stack of fully connected layers was used between the stack of convolution layers and the scalar output node trained to contribute to similarity estimation.

For DESK, the most successful architectures were *Siamese* (Chopra et al., 2005), where the two input images were processed using two paired stacks of convolutional layers that had disjoint connections but shared (identical) weights, followed by a common stack of fully connected layers. Other architectures that we tried but found to be not as well-suited for estimating image pair similarity are shown in Figure 2. For other architectures that might work, the reader is referred to (Zagoruyko & Komodakis, 2015).

3.2. Levels of transfer learning

To address our first goal of practical efficiency in training a kernel, we decided to not train the convolutional layers of the Siamese net from scratch at all. We simply picked pre-trained CNNs trained to classify an appropriate set of non-target classes and duplicated their convolutional layers in the two paired stacks of convolutional layers in the Siamese net. We did not fine-tune these layers at all after that. This level of transfer learning was common to all of our experiments. We experimented with two transfer learning schemes or levels:

1. **Conv-transfer:** In the first scheme, after copying the convolutional layers from non-target classes and freezing them, we trained the shared stack of fully connected layers on pairs of images (processed through the convolutional layers) from *target* classes. This was done by either fine-tuning fully connected

layers pre-trained on image pairs from non-target classes by using a image pairs from target classes, or training from scratch on target classes.

2. **Full-transfer:** In the second scheme, in addition to copying convolutional layers from non-target classes and freezing them, we also trained the shared stack of fully connected layers on pairs of images (processed through the paired conv layers) from *non-target* classes. Then we froze the entire Siamese net while processing the target classes.

Interestingly, even the conv-transfer learning scheme requires only a small number of samples from target classes to train the fully connected layers because a small number of images can generate a large number of image *pairs*.

3.3. Supervised training of Siamese net

The notion of similarity that we propose is the probability of an image pair to belong to the same class no matter what their classes are. The main reason for choosing this notion of similarity as opposed to, say, semantic distance between class labels words, was to facilitate comparison with CNNs that use the same supervised information. We wanted to test the hypothesis that *learning to compare is inherently a better objective than learning to classify in terms of training sample and time efficiency.*

We trained the shared fully connected layers of the Siamese nets with CNN features for a pair of images computed using frozen conv layers as inputs. Supervised target output of $\{0, 1\}$ was emulated during training, where 1 was used to code similar image pairs (belonging to the same class), and 0 for dissimilar pairs (belonging to different classes). We used an equal number of similar and dissimilar pairs for training.

3.4. From Siamese net to kernel

According to Mercer's theorem, a function of two inputs represents a reproducing kernel Hilbert space if it is (a) symmetric, and (b) positive semi-definite (PSD) (Schölkopf & Smola, 2002). The Siamese net output in DESK is not guaranteed to be that function, which can occasionally cause convergence problems for SVM packages training on gram matrices that are not PSD. Therefore, before training an SVM, which requires such a kernel for *guaranteed* convex objective function, we experimented with the following post-processing schemes for the gram matrix S with elements s_{ij}:

1. **Co-incidence matrix:** It has been suggested that when a non-negative similarity metric captured in gram matrix S is asymmetric, it can be converted to a co-incidence matrix $S^T S$ that is PSD and can also be used for classification in an SVM (Schölkopf et al., 2001). The added advantage of this scheme in processing unseen target classes is that computing co-incidence confers another level of learning where each sample's vector of distances with other samples is compared to that of other vectors. So, it transforms S into a matrix $S^T S$ that really is tailored to the inter-distances from the samples that generated of S itself. In our experiments, this scheme gave consistently high separation of test cases and high classification performance.

2. **Pair exchange and ignoring negative eigen-values:** First we exchanged the inputs and took the average of the two kernel computations, that is, we used the matrix $\frac{1}{2}(S + S^T)$. Then we computed eigen decomposition of this symmetric matrix, set the negative eigenvalues to zero, and back-projected the eigenvectors using the new eigenvalues to obtain a symmetric and PSD matrix. We used this matrix for training the SVM, and used the projection of the symmetrized testing gram matrix into this eigenspace for testing the SVM. This scheme didn't perform as well as the co-incidence scheme.

3. **Pick-out kernel:** Munoz et. *al.* presented a method, known as *pick-out*, that specifically takes classification labels into account to build the proximity matrix for enforcing positive semidefiniteness on the asymmetric Gram matrices (Alberto Munoz & Moguerza, 2003). In this method, while training, the $max(s_{ij}, s_{ji})$ is used as the kernel value if both samples (images) i and j belong to the same class, and $min(s_{ij}, s_{ji})$ otherwise. For testing, one first assumes that the test sample belongs to one class to compute its distance from the separating hyperplane, and then repeats the exercise with assuming that the test belongs to the other class. The binary class decision is taken based on comparison of the two signed distances. Its extension to more than two classes isn't obvious, and it also gave worse performance than co-incidence matrix for Dogs vs. Cats binary classification.

3.5. SVM training

For obvious reasons, the SVM was always trained using kernel matrices from target classes. We tested its performance on kernels derived from conv-transfer and full-transfer schemes for different number of training samples. For conv-transfer, we used the same samples to train the fully-connected layers of the Siamese net and the SVM. SVM training took around 5 to 10 minutes.

4. Experiments and Results

In this Section, we describe the data used for target classes, and their non-target classes as well as the results of our experiments.

4.1. Data sets and DESK architectures

We probed questions around the properties of the learned kernel as well as the recognition performance when the kernel was used in SVM. To do so, we selected three datasets of target classes along with their respective non-target classes. The non-target classes were selected such that the images were of the same type as those of the target classes in terms of size and broad categories. For example, two hand-written scripts share pen strokes, while two sets of natural images share similar features and feature hierarchy. This pairing of target and non-target classes is shown in Table 1.

Table 1. Target and non-target classes

Target classes	#	Non-target classes
Dogs vs. Cats (Dogs vs. Cats)	2	ImageNet sans dogs and cats (Deng et al., 2009)
Handwritten digits (MNIST) (MNIST)	10	Handwritten alphabet (NIST) (NIST)
CIFAR-10 objects (Alex Krizhevsky & Hinton)	10	CIFAR-100 objects (Alex Krizhevsky & Hinton)

We selected cat vs. dog classification task because it is binary and gives better insights into the performance of the kernel and the SVM without having to interpret how these scale to multi-class problems. It should be noted that this is not a trivial task because both dogs and cats are furry mammals usually pictured in similar surroundings. While the other two target datasets and their non-target counterparts had a fixed image size, ImageNet had variable image sizes. We standardized the images in ImageNet and to $227\times227\times3$ by scaling.

The best performing architectures for the two of the datasets used in our experiments as defined by kernel accuracy are shown in Table 2, where FC represents fully connected layers whose number of neurons are mentioned, and Conv represents Siamese convolutional and pooling layers whose filter sizes, number of filters, and pooling sizes respectively are mentioned. Note that all nonlinearities were rectified linear units (ReLU), except for the output node, which had a logistic sigmoid nonlinearity. We used AlexNet pre-trained architecture for ImageNet (A. Krizhevsky et. al., 2012). For CIFAR dataset, we used mxnet architecture (MXNet), but we do not show it here because it is highly complex. The convolutional layers for NIST were trained from scratch using a CNN that gave approximation 92% classification accuracy on English alphabet. The FC layers were trained using hinge loss cost function as suggested in (Zagoruyko & Komodakis, 2015) and used a dropout of 0.5.

Table 2. Neural network architectures. In the convolutional layers, filter size, number of filters, and pooling size are mentioned respectively.

Data set	ImageNet	NIST
FC 3	1,000	
FC 2	4,096	
FC 1	4,096	800
Conv 5	3×3, 256, 3×3	
Conv 4	3×3, 384, 1×1	
Conv 3	3×3, 384, 1×1	3×3, 128, 2×2
Conv 2	5×5, 256, 3×3	3×3, 64, 2×2
Conv 1	11×11, 96, 3×3	3×3, 32, 2×2
Image	227×227×3	32×32×1

4.2. Kernel performance

We tested the kernel's performance for both transfer schemes described in Section 3.2. While conv-transfer was our main focus, full-transfer shows the generalization capabilities of features estimated for learning to compare as opposed to learning to classify.

4.2.1. Conv-transfer

Our main experiments were about only transferring the convolutional layers. Then, we trained the fully connected layers on training images (actually, pairs thereof) from the target classes using 500, 1,000 and 5,000 images. This represents the conv-transfer learning scheme. However, as shown in Table 3, we used a lot more pairs of images that we could generate from the paltry number of images. We then tested the kernel on 10,000 pairs from a held-out set of 1,000 images from the target classes. We generated an equal number of similar and dissimilar pairs for both training and testing.

The kernel accuracy decreased as the number of training images was increased. However, this lower kernel performance due to more extensive training need not necessarily lead to worse recognition performance, as the kernel was trained to find a more meaningful notion of similarity on a larger set of images. Kernel accuracy can give a crude approximation of 1-shot learning performance for a binary classification problem.

4.2.2. Full-transfer

To test generalization capabilities of the kernel, we tested the first transfer learning scheme described in Section 3.2

Table 3. Kernel testing AUC for training fully connected layers on the different number of samples from target classes.

Images	Pairs	Dogs Cats	MNIST	CIFAR-10
500	10,000	0.999	0.996	0.996
1,000	60,000	0.999	0.994	0.999
5,000	60,000	0.999	0.993	0.997

by training the fully connected layers on pairs of CNN features from non-target classes. We used only 5,000 images but 60,000 pairs to train the kernel. We then tested the kernel on 10,000 image pairs from 1,000 images of unseen target classes, and the results are reported in Table 4. The performance was surprisingly high for binary classification, and still encouraging for the 10 class sets. Among the latter, digit comparison was better generalized based on learning to compare alphabet. This was expected because generalization in comparing pen strokes of scripts is inherently easier than features of natural objects.

The kernel took only between 1 to 3 hours to train on a hexa-core 16GB RAM machine with a 2000 CUDA®core GPU.

Table 4. Kernel generalization AUC on target (test) classes for training all layers on non-target classes.

Images	Pairs	Dogs Cats	MNIST	CIFAR-10
5,000	60,000	0.990	0.943	0.868

4.3. Kernel visualization

We used visualization of class separation to gain more insights into the relative performance of DESK and its alternative, which is to use CNN features. Tools such as t-SNE are available to visualize class separation for high dimensional data, by projecting the data into two dimensions where their pair-wise distances in small neighborhoods are representative of their distances in the original space (Maaten & Hinton, 2008). One can also establish a correspondence between explicit features of a CNN and implicit features of a kernel by using columns of the kernel's gram matrix in lieu of features.

We tried to gain insight into any inherent advantage of DESK over CNN, impact of fine-tuning, and the advantage of using the co-incidence matrix derived from the output of the Siamese net using post-processing. For this, we visualized CNN features, Siamese net output, and DESK output (Siamese with post-processing) for both full-transfer (without fine-tuning) and conv-transfer (with fine-tuning). These results are shown in Figures 3 and 4, and are very striking. DESK seems to do a much better job at separating the classes compared to CNN even before fine-tuning. After fine-tuning, the results are remarkably well-separated with only small clusters of confusion.

4.4. Classification performance using DESK in an SVM

We now report our main result, which is classification accuracy using a small number of training samples from target classes.

4.4.1. Conv-transfer

For the conv-transfer learning scheme where only the convolutional layers were trained on the non-target classes, we got some very encouraging recognition results. After freezing the convolutional layers, we trained the fully connected layers and the SVM on image pairs generated from only 500, 1,000 or 5,000 training images from target classes. This is very frugal considering that most reported work on NIST and CIFAR datasets used around 50,000 training images. For 500 samples, we generated 10,000 image pairs, and for the other two cases we used 60,000 pairs (half of them from the matched classes) for training DESK. Note that the SVM had a slack penalty as a hyper-parameter, which was fixed for testing using a validation subset taken from the training set. We then tested the SVM on 5,000 test images from the target classes. We compared these results with previously reported classification results for the test data sets as well as some obvious alternative transfer learning schemes that use a small number of training examples from target classes. These included use of pre-trained CNN features on non-target classes passed to an SVM via a static kernels (linear and RBF) (Huang & LeCun, 2006b), or coupling pre-trained CNN features with new fully connected layers and fine-tuning on a small number of examples from target classes (Alexey Dosovitskiy, 2015). DESK, which essentially *learns* a kernel on top of CNN features, generally outperformed these methods in 500 to 5,000 sample range as shown in Table 5. It was observed that the performance of fine-tuned CNN goes up faster with number of training samples than that of DESK, possibly because CNN learns class-appearance-specific weights, while DESK learns comparison-specific weights. Pre-training in our experiments was done using a large number of labeled samples from the mentioned non-target classes.

These results are quite encouraging in the sense that with as few as $1,000$ samples, the recognition rates are close to state of the art such as 0.998 for MNIST (Wan et al., 2013) and 0.84 for CIFAR-10 (Hinton et al., 2012) using CNNs (with ReLU and dropout) that were trained for classification on approximately $50,000$ samples. These results outperform previously published results for a similar number of samples on these datasets as well as other transfer learning techniques that we implemented.

By comparing Tables 3 and 5 it is clear that although the kernel accuracy decreased when the number of training im-

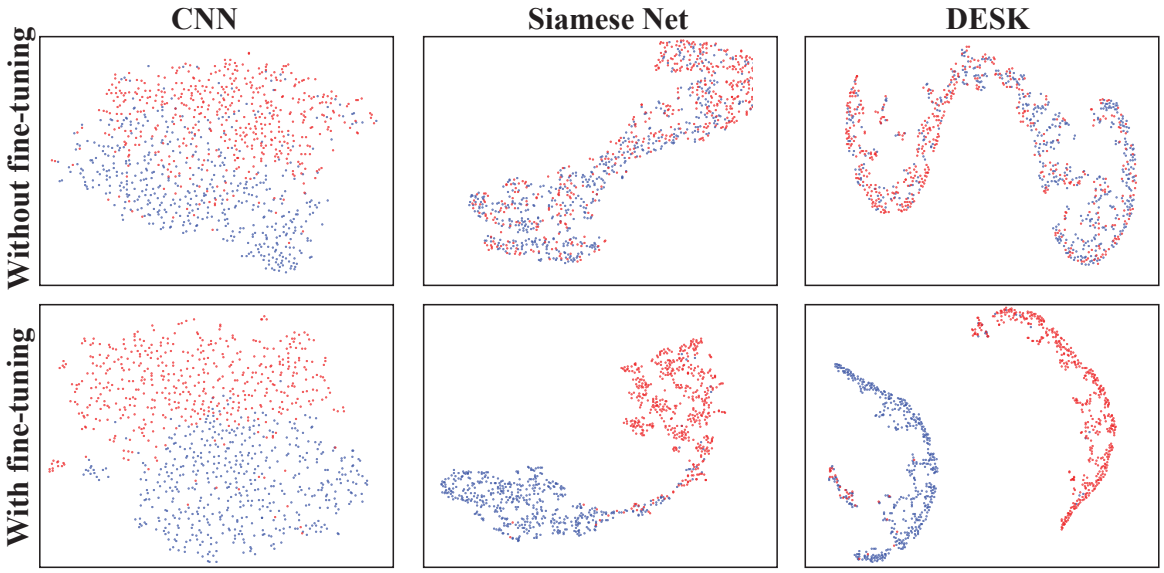

Figure 3. Class separation of Dogs vs. Cats for CNN features, Siamese net, and DESK (Siamese net with post-processing) for the two transfer learning schemes (all plots are for 1000 testing samples of target classes).

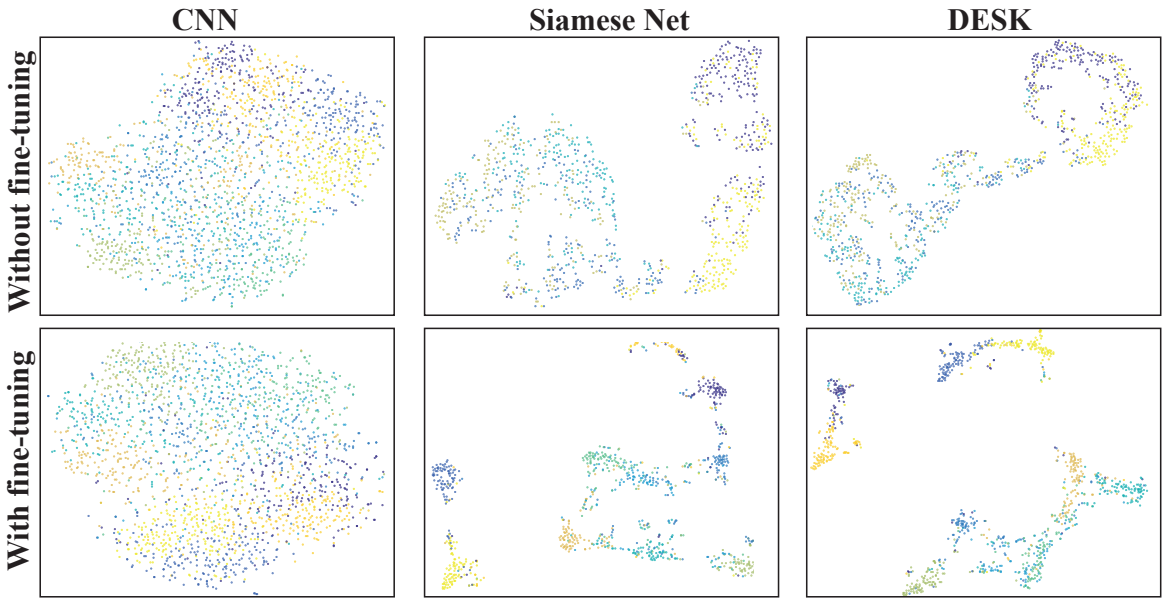

Figure 4. Class separation of CIFAR-10 object classes for CNN features, Siamese net, and DESK (Siamese net with post-processing) for the two transfer learning schemes (all plots are for 2000 testing samples of target classes).

Table 5. Classification accuracy vs. number of training samples from target classes for DESK and its alternatives, including the best reported results for MNIST(Philip Bachman, 2014) and CIFAR-10(Alexey Dosovitskiy, 2015).

Dataset	Dogs Cats			MNIST			CIFAR-10		
# Training samples from target classes	500	1,000	5,000	500	1,000	5,000	500	1,000	5,000
Previously published best results	None	None	None	.976	.978	.981	None	.774 at 4,000	
CNN features + linear kernel + SVM	.963	.965	.967	.937	.946	.971	.600	.628	.744
CNN features + RBF kernel + SVM	.963	.968	.973	.926	.956	.978	.661	.705	.758
CNN features + fine-tuning	.810	.890	.976	.797	.881	.952	.617	.721	**.840**
DESK + SVM	**.996**	**.995**	**.997**	**.989**	**.995**	**.993**	**.758**	**.774**	.809

ages (not pairs) was increased, the classification accuracy either increased or stayed about the same. Thus, although the kernel performed worse when trained on pairs taken from a larger set of images, it perhaps learned a better representation of the data and task, thus aiding the SVM in testing performance.

4.4.2. FULL-TRANSFER

We next report classification accuracy using a full-transfer kernel in Table 6. That is, even the fully connected layers of the kernel were trained using pairs from non-target classes. Using this transferred kernel, only the SVM was trained on target classes. For dogs vs. cats and digits, the performance is remarkably high considering that only training the SVM on the target classes. This, along with results of Table 4, show that a general image similarity can be learned without even fine-tuning the weights of DESK with only using class labels as information.

Table 6. SVM classification testing accuracy when the entire kernel was transferred from non-target classes.

Images	Pairs	Dogs Cats	MNIST	CIFAR-10
5,000	60,000	0.987	0.940	0.607

5. Discussion and Conclusions

In this paper, we showed that a Siamese deep neural network architecture that takes two images and estimates their similarity can be used in a trainable kernel – DESK. We showed that such a network can be trained to capture similarity between two images using supervised learning, where such similarity can be defined by an image pair belonging to the same class, no matter what that class is. With specific modifications to the resultant gram matrix, it can be used in an SVM for classification.

DESK lends itself to transfer learning from non-target classes. With a few hundred to a few thousand training images from target classes, this kernel can be trained to give classification accuracy competitive with traditional CNNs trained on tens of thousands training images to predict target classes. That is, it generalizes to estimating similarity of completely unseen classes, or with only a small amount of training data from the target classes. Consequently, its training time is quite less as well; hours, instead of days. When the entire kernel is pre-trained wholly on tens of thousands of samples from non-target classes, then only the SVM needs to be trained on target classes, which takes only a few minutes. Thus, its main advantage is to be able to train a classifier with far fewer samples and training time on target classes than traditional CNNs in cases when a lot of labeled data is not available.

The main insight from this study is that learning to compare image pairs generalizes better with far fewer samples of target classes than learning to classify. This might be due to offloading the need to code appearance to support vectors instead of storing appearance code templates in the weights of the higher layers in a Siamese architecture. Surprisingly, this can be done using just class labels as information from non-target classes. Using richer semantic information than class labels may lead to even better kernels.

The main disadvantage of DESK is its testing time, as the kernel needs to be computed as a gram matrix of all pairs of testing and training samples, although internally SVM packages only choose to use the columns corresponding to the training support vectors. This load can be reduced if we compute the kernel on pairs of testing samples and only those training samples that correspond to support vectors, which is not allowed in most SVM packages, although it is easy to implement. Moreover, the convolutional part of both support vectors and test samples can be computed only once (separately for each image instead of all pairs of images) due to the disjoint (although Siamese) architecture up to the fully connected layers. Then, only the outputs of fully connected layers need to be computed on all image pairs.

Ways to ensure symmetry of output with respect to the inputs or its positive semi-definiteness right out of the network will be useful to eliminate some of the post-processing that we used. Different variations of SVM and CNN can also be tried to optimize the performance. For example, it may be possible to improve the performance by training the neural network in a deeply supervised framework (Lee et al., 2014). The meaningfulness of the kernel can perhaps be further improved for object classification on datasets like CIFAR-10 by not including those pairs from CIFAR-100 that belong to related classes. This is because classes such as dogs and wolves are neither similar nor dissimilar. That is, they are neither the same, nor as far apart as an airplane and a dog. This is where a more granular notion of semantics such as word relational hierarchies may come in handy.

References

A. Krizhevsky et. al. Imagenet classification with deep convolutional neural networks. In *International Conference on Neural Information Processing Systems (NIPS)*, 2012.

Alberto Munoz, Isaac Martin de Diego and Moguerza, Javier M. Support vector machine classifiers for asymmetric proximities. In *International Conference on Artificial Neural Networks*, 2003.

Alex Krizhevsky, Vinod Nair and Hinton, Geoffrey. Cifar Dataset https://www.cs.toronto.edu/

~kriz/cifar.html. doi: http://www.cs.toronto.edu/~kriz/cifar.html.

Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg Martin Riedmiller Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2015.

Boureau, Y-Lan, Ponce, Jean, and Lecun, Yann. A theoretical analysis of feature pooling in visual recognition. In *27th International Conference on Machine Learning (ICML)*, 2010.

Caffe. Berkeley Vision and Learning Center - Caffe http://caffe.berkeleyvision.org/. doi: \url{http://caffe.berkeleyvision.org/}.

Chapelle, O., Haffner, P., and Vapnik, V. N. Support vector machines for histogram-based image classification. *Transactions on Neural Networks*, 10(5):1055–1064, September 1999. ISSN 1045-9227. doi: 10.1109/72.788646. URL http://dx.doi.org/10.1109/72.788646.

Cho, Youngmin and Saul, Lawrence K. Kernel methods for deep learning. In Bengio, Y., Schuurmans, D., Lafferty, J.D., Williams, C.K.I., and Culotta, A. (eds.), *Advances in Neural Information Processing Systems 22*, pp. 342–350. Curran Associates, Inc., 2009. URL http://papers.nips.cc/paper/3628-kernel-methods-for-deep-learning.pdf.

Chopra, S., Hadsell, R., and LeCun, Y. Learning a similarity metric discriminatively, with application to face verification. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pp. 539–546 vol. 1, June 2005. doi: 10.1109/CVPR.2005.202.

Dahl, George E, Sainath, Tara N, and Hinton, Geoffrey E. Improving deep neural networks for lvcsr using rectified linear units and dropout. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pp. 8609–8613. IEEE, 2013.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.

Dogs vs. Cats. Kaggle.com - Dogs vs. Cats https://www.kaggle.com/c/dogs-vs-cats/data. doi: https://github.com/dmlc/mxnet.

Gönen, Mehmet and Alpaydin, Ethem. Multiple kernel learning algorithms. *Journal of Machine Learning Research*, 12:2211–2268, July 2011. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=1953048.2021071.

Hinton, Geoffrey E., Srivastava, Nitish, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Improving neural networks by preventing co-adaptation of feature detectors. *CoRR*, abs/1207.0580, 2012. URL http://dblp.uni-trier.de/db/journals/corr/corr1207.html#abs-1207-0580.

Huang, Fu Jie and LeCun, Y. Large-scale learning with svm and convolutional for generic object categorization. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 1, pp. 284–291, 2006a. doi: 10.1109/CVPR.2006.164.

Huang, Fu Jie and LeCun, Yann. Large-scale learning with svm and convolutional for generic object categorization. In *Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Volume 1*, CVPR '06, pp. 284–291, Washington, DC, USA, 2006b. IEEE Computer Society. ISBN 0-7695-2597-0. doi: 10.1109/CVPR.2006.164. URL http://dx.doi.org/10.1109/CVPR.2006.164.

Junlin, Hu, Jiwen, Lu, and Tan, Yap-Peng. Discriminative deep metric learning for face verification in the wild. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pp. 1875–1882. IEEE, 2014.

Lanckriet, Gert R. G., Cristianini, Nello, Bartlett, Peter, Ghaoui, Laurent El, and Jordan, Michael I. Learning the kernel matrix with semidefinite programming. *J. Mach. Learn. Res.*, 5:27–72, December 2004. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=1005332.1005334.

Lee, Chen-Yu, Xie, Saining, Gallagher, Patrick, Zhang, Zhengyou, and Tu, Zhuowen. Deeply-supervised nets. *arXiv preprint arXiv:1409.5185*, 2014.

LeNet. LeNet http://yann.lecun.com/exdb/lenet/ . doi: http://yann.lecun.com/exdb/lenet/.

Maaten and Hinton, Geoffrey E. Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

Mairal, Julien, Koniusz, Piotr, Harchaoui, Zaïd, and Schmid, Cordelia. Convolutional kernel networks. *CoRR*, abs/1406.3332, 2014. URL http://arxiv.org/abs/1406.3332.

MNIST. MNIST Dataset http://yann.lecun.com/exdb/mnist/. doi: http://yann.lecun.com/exdb/mnist/.

Montavon, Grégoire, Braun, Mikio L, and Müller, Klaus-Robert. Kernel analysis of deep networks. *The Journal of Machine Learning Research*, 12:2563–2581, 2011.

MXNet. MXNet GitHub https://github.com/dmlc/mxnet. doi: https://github.com/dmlc/mxnet.

NIST. NIST Handprinted Forms and Characters Database http://www.nist.gov/srd/nistsd19.cfm. doi: http://www.nist.gov/srd/nistsd19.cfm.

Philip Bachman, Ouais Alsharif, Doina Precup. Learning with pseudo-ensembles. In *Neural Information Processing Systems (NIPS) Conference*, 2014.

Schölkopf, Bernhard and Smola, Alexander J. *Learning with kernels: Support vector machines, regularization, optimization, and beyond*. MIT press, 2002.

Schölkopf, Bernhard, Herbrich, Ralf, and Smola, Alex J. A generalized representer theorem. In *Proceedings of the 14th Annual Conference on Computational Learning Theory and and 5th European Conference on Computational Learning Theory*, COLT '01/EuroCOLT '01, pp. 416–426, London, UK, UK, 2001. Springer-Verlag. ISBN 3-540-42343-5. URL http://dl.acm.org/citation.cfm?id=648300.755324.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.

Sonnenburg, Sören, Rätsch, Gunnar, Schäfer, Christin, and Schölkopf, Bernhard. Large scale multiple kernel learning. *Journal of Machine Learning Research*, 7:1531–1565, December 2006. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=1248547.1248604.

Srivastava, Nitish, Hinton, Geoffrey, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014. URL http://jmlr.org/papers/v15/srivastava14a.html.

Theano. DeepLearning.net - Theano http://deeplearning.net/software/theano/. doi: http://deeplearning.net/software/theano/.

Wan, Li, Zeiler, Matthew D., Zhang, Sixin, LeCun, Yann, and Fergus, Rob. Regularization of neural networks using dropconnect. In *ICML (3)*, volume 28 of *JMLR Proceedings*, pp. 1058–1066. JMLR.org, 2013. URL http://dblp.uni-trier.de/db/conf/icml/icml2013.html#WanZZLF13.

Yosinski, Jason, Clune, Jeff, Bengio, Yoshua, and Lipson, Hod. How transferable are features in deep neural networks? *CoRR*, abs/1411.1792, 2014. URL http://arxiv.org/abs/1411.1792.

Zagoruyko, Sergey and Komodakis, Nikos. Learning to compare image patches via convolutional neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.

Zou, Will Y, Ng, Andrew Y, and Yu, Kai. Unsupervised learning of visual invariance with temporal coherence. In *NIPS 2011 Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.